

REMARKS

Claims 1-45 are currently pending in the application. By this response, no claims are amended, added, or canceled. Reconsideration of the rejected claims in view of the above amendments and the following remarks is respectfully requested.

Interview Summary

Applicants appreciate the courtesy extended by Examiner Mirza in granting and conducting a telephone interview with Applicants' representative on July 17, 2009 ("Interview"). In the Interview, Applicants' representative pointed out that the Bruckert reference (U.S. Pub. No. 2002/0049859) does not appear to teach the features of the claimed invention as asserted by the Examiner in the Office Action. More specifically, Applicants' representative discussed how Bruckert does not appear to teach anything related to burden of a server or load-balancing, and therefore, does not appear to teach the recited steps that are based on a server being burdened and unburdened. The Examiner did not disagree, and asked that the arguments be made in a formal written response. The Examiner agreed to consider the arguments and issue a new non-final action, if warranted. However, agreement with respect to allowable claim language was not reached.

35 U.S.C. §103 Rejection

Claims 1-45 are rejected under 35 U.S.C. §103(a) for being unpatentable over U.S. Pat. No. 7,213,065 ("Watt"), in view of U.S. Pub. No. 2002/0049859 ("Bruckert"). This rejection is respectfully traversed.

To establish a *prima facie* case of obviousness, all claim limitations must be taught or suggested by the prior art. *See, In re Royka*, 490 F.2d 981, 985, 180 USPQ 580, 583 (CCPA 1974); *see also, In re Vaeck*, 947 F.2d 488, 20 USPQ2d 1438 (Fed. Cir. 1991).¹ If the prior art reference(s) do not teach or suggest all of the claim limitations, Office personnel must explain why the differences between the prior art and the claimed invention would have been obvious to one of ordinary skill in the art (MPEP 2141). Applicants submit that no proper combination of the applied art teaches or suggests each and every feature of the claimed invention.

Initially, Applicants submit that the rejection of at least independent claims 1 and 24 is improper because the Examiner has not addressed the language of these claims in the Office Action. Instead, the Examiner has grouped claims 1 and 24 with the rejection of claim 18 while only addressing the language of claim 18. However, claims 1 and 24 recite features that are not recited in claim 18, such that an explanation of the rejection of claim 18 cannot serve to establish a *prima facie* case of obviousness with respect to claims 1 and 24.

For example, claims 1 and 24 each recite *defining a plurality of virtual clusters from a plurality of servers, and routing a request to a virtual cluster of the plurality of virtual clusters based on predetermined criteria in order to allocate system resources*. The Examiner has not identified any teaching of these features in the applied art. In fact, the Examiner did not even address these features in the explanation of the rejection. MPEP §707.07(d), states that “[a] plurality of claims should never be grouped together in a common rejection, unless that rejection is equally applicable to all claims in the group.” In this case, the Examiner has improperly grouped claims 1 and 24 with the rejection of other claims while failing to address the language

¹ While the *KSR* court rejected a rigid application of the teaching, suggestion, or motivation (“TSM”) test in an obviousness inquiry, the [Supreme] Court acknowledged the importance of identifying “a reason that would have prompted a person of ordinary skill in the relevant field to combine the elements in the way the claimed new invention does” in an obviousness determination. *Takeda Chemical Industries, Ltd. v. Alphapharm Pty., Ltd.*, 492

of claims 1 and 24. Therefore, the rejection of claims 1 and 24 is improper and should be withdrawn for at least this reason.

In any event, Applicants submit that the applied art does not disclose or suggest all of the features of the claimed invention. The invention generally relates a system and method of monitoring servers in a network and, more particularly, intelligently routing requests from a network dispatcher to servers allocated in virtual clusters based on rules. More specifically, independent claim 1 recites:

1. A method of processing queries in a network, comprising the steps of:
 - defining a plurality of virtual clusters from a plurality of servers;
 - routing a request to a virtual cluster of the plurality of virtual clusters based on predetermined criteria in order to allocate system resources;
 - removing at least one of the plurality of servers from the virtual cluster when at least one of the plurality of servers is burdened;
 - creating a new virtual cluster comprising only the removed at least one of the plurality of servers; and
 - returning the removed at least one of the plurality of servers back to the virtual cluster when the at least one of the plurality of servers is unburdened.

Also, independent claim 18 recites:

18. A method for load balancing servers, comprising the steps of:
 - allocating a plurality of servers among a plurality of virtual clusters;
 - monitoring the plurality of virtual clusters for workload capacity;
 - removing at least one of the plurality of servers from the plurality of virtual clusters when at least one of the plurality of servers is burdened;
 - creating a new virtual cluster comprising only the removed at least one of the plurality of servers;
 - returning the removed at least one of the plurality of servers back to the plurality of virtual clusters when the at least one of the plurality of servers is unburdened; and

F.3d 1350, 1356-1357 (Fed. Cir. 2007) (quoting *KSR International Co. v. Teleflex Inc.*, 127 S.Ct. 1727, 1731 (2007)).

reassigning at least one server from one of the plurality of virtual clusters to another of the plurality of virtual clusters based on workload capacity of the at least one server in order to reallocate system resources.

Additionally, independent claim 24 recites:

24. A computer program product comprising a computer usable medium having readable program code embodied in the medium, the computer program product includes at least one component to:
define a plurality of virtual clusters from a plurality of servers;
route a request to a virtual cluster of the plurality of virtual clusters based on predetermined criteria to allocate system resources;
remove at least one of the plurality of servers from the virtual cluster when at least one of the plurality of servers is burdened;
create a new virtual cluster comprising only the removed at least one of the plurality of servers; and
return the removed at least one of the plurality of servers back to the virtual cluster when the at least one of the plurality of servers is unburdened.

The Examiner asserts that Watt discloses allocating a plurality of servers among a plurality of virtual clusters, and monitoring the plurality of virtual clusters for workload capacity at lines 36-56 of col. 4. The Examiner also asserts that Watt discloses reassigning at least one server from one of the plurality of virtual clusters to another of the plurality of virtual clusters based on workload capacity of the at least one server in order to reallocate system resources at lines 34-47 of col. 2.

The Examiner acknowledges that Watt does not teach removing at least one of the plurality of servers from the virtual cluster when at least one of the plurality of servers is burdened, creating a new virtual cluster comprising only the removed at least one of the plurality of servers, or returning the removed at least one of the plurality of servers back to the virtual cluster when the at least one of the plurality of servers is unburdened. Applicants agree that Watt does not disclose or suggest these features.

The Examiner asserts that Bruckert teaches, at paragraph 0040, the above-noted features of the claimed invention that are missing from Watt. The Examiner concludes that it would have been obvious to modify Watt in view of the teachings of Bruckert, and the combination of teachings renders the claimed invention unpatentable. Applicants disagree with the conclusion of obviousness for the following reasons.

Contrary to the Examiner's assertions, Bruckert does not disclose or suggest: *removing at least one of the plurality of servers from the plurality of virtual clusters when at least one of the plurality of servers is burdened; creating a new virtual cluster comprising only the removed at least one of the plurality of servers; or, returning the removed at least one of the plurality of servers back to the plurality of virtual clusters when the at least one of the plurality of servers is unburdened*, as recited in the claimed invention. As discussed in the above-noted Interview, Bruckert makes no mention of server burden/unburden or of load-balancing. Instead, Bruckert is concerned with scaling clusters (adding and/or removing nodes to a cluster) and maintaining operability of the clusters during scaling (paragraph 0007, lines 1-4; paragraph 0038, lines 1-11). Bruckert's invention involves a node-naming-agent (NNA) that provides support for scaled clustering by transforming a local/global cluster address into a corresponding global/local cluster address for each packet in an outbound/inbound path. Bruckert's primary object is providing naming (e.g., addressing) that permits scalability that is transparent to the applications using the clusters.

Particularly, at paragraphs 0039 and 0040, Bruckert describes two prior art naming (e.g., addressing) techniques that can be used when scaling a cluster. In paragraph 0039, Bruckert describes that each and every node that may be used in the clustered system is assigned a unique address referred to as a cluster ID. In this manner, when connecting device together, you never

have a problem of duplicate IDs, and the cluster can provide continuous service without having downtime for re-assigning addresses.

In paragraph 0040, Bruckert describes an alternative method where, instead of assigning a unique ID to each and every device, only devices actually added to the cluster are assigned an ID. In other words, the same cluster IDs may be re-assigned to different devices based on which devices are added to the cluster. Bruckert contrasts this second approach (e.g., paragraph 0040) with the previously described approach (e.g., paragraph 0039) by noting that re-assigning of IDs impacts network traffic because a node that is being re-assigned will be temporarily unavailable until its ID is assigned.

However, Bruckert makes no mention of burden on a server, or of load balancing. Instead, Bruckert appears to be solely interested in the methods of performing the actual scaling (adding and/or removing nodes from a cluster), to make operation of clusters continuous without any suspension of operation during such scaling. Because Bruckert does not disclose anything to do with server burden, Bruckert cannot reasonably be construed as teaching *removing at least one of the plurality of servers from the plurality of virtual clusters when at least one of the plurality of servers is burdened*, as recited in the claimed invention. Moreover, since Bruckert is completely silent with respect to burdening and unburdening of a server, Bruckert does not teach *returning the removed at least one of the plurality of servers back to the plurality of virtual clusters when the at least one of the plurality of servers is unburdened*. Therefore, Bruckert fails to compensate for the admitted deficiencies of Watt with respect to these features. Accordingly, the applied art does not teach or suggest all of the features of the claimed invention.

Furthermore, neither Watt nor Bruckert teaches *creating a new virtual cluster comprising only the removed at least one of the plurality of servers*, as recited in the claimed invention. The

Examiner admits, and Applicants agree, that Watt does not teach this feature. Moreover, the Examiner has failed to identify this feature in Bruckert. Instead, paragraph 0040 of Bruckert, which is cited by the Examiner in the Office Action, merely teaches re-assigning cluster IDs when nodes are joined to a cluster. This passage of Bruckert makes no mention of creating a new virtual cluster comprising only a server that was removed from another cluster. Therefore, the applied art fails to teach *creating a new virtual cluster comprising only the removed at least one of the plurality of servers*, as recited in the claimed invention.

Moreover, even assuming for argument sake that Bruckert can be construed as teaching *creating a new virtual cluster comprising only the removed at least one of the plurality of servers*, Applicants submit that it would not have been obvious to add this feature to Watt because this would change the principle of operation of the Watt method and system. According to MPEP 2143.01:

If the proposed modification or combination of the prior art would change the principle of operation of the prior art invention being modified, then the teachings of the references are not sufficient to render the claims *prima facie* obvious. *In re Ratti*, 270 F.2d 810, 123 USPQ 349 (CCPA 1959) ... The court reversed the rejection holding the "suggested combination of references would require a substantial reconstruction and redesign of the elements shown in [the primary reference] as well as a change in the basic principle under which the [primary reference] construction was designed to operate." 270 F.2d at 813, 123 USPQ at 352.).

Watt discloses adding a server from a "free pool" when a server pool is overloaded, and returning a server to the free pool when the server pool is underloaded. Watt makes no mention of creating a new virtual cluster comprising only a removed server. Instead, Watt moves servers into and out of the free pool, which is an existing pool (not a new virtual cluster). If one were to modify Watt to create a new virtual cluster comprising only a removed server, then the principle

of operation would be changed. Therefore, according to MPEP 2143.01, such a modification would not have been obvious.

For all of the above-noted reasons, Applicants submit that the applied art does not render independent claims 1, 18 and 24 unpatentable. Claims 2-17, 19-23, and 25-45 depend from allowable independent claims 1, 18, and 24, respectively, and are allowable at least for the same reasons as the respective base claims. Moreover, the applied art fails to teach many of the features recited in the dependent claims.

For example, claim 41 depends from allowable independent claim 1, and additionally recites *projecting a rate of routing to each of the plurality of servers*. The Examiner asserts that Bruckert teaches this feature at paragraph 0010. Applicants disagree. Paragraph 0010 of Bruckert is reproduced below.

[0010] Further in accordance with the purpose of the invention, as embodied and broadly described herein, the invention relates to a method including steps for scaling the clustered system. Additionally, the invention relates to a computer readable medium in a scalable clustered system that embodies computer program code configured to cause that system to perform steps for configuring and scaling that system. The steps include operatively linking two or more cluster nodes via a global fabric in order to form a larger clustered system. Each of the cluster nodes has end nodes and a local fabric interconnecting the end nodes. The steps further include routing global packet traffic between the two or more cluster nodes in the larger clustered system via the global fabric; and routing local packet traffic between the one or more end nodes in each of the cluster nodes via the local fabric. The steps additionally include operatively interposing an NNA between the local fabric and the global fabric.

There is nothing in this passage that teaches projecting a rate of routing to each of the plurality of servers. Applicants acknowledge that Bruckert discloses routing packets between cluster nodes. However, Bruckert makes no mention of projecting anything, and makes no mention of a rate of routing. Therefore, the applied art does not teach *projecting a rate of routing to each of the plurality of servers*, as recited in claim 41.

Claim 42 depends from claim 41 and additionally recites *adjusting the rate of routing based on a relative degree of overload on at least one of the plurality of servers*. As with claim 41, the Examiner asserts that Bruckert teaches this feature at paragraph 0010. Applicants disagree.

Bruckert makes no mention of a rate of routing in paragraph 0010. Instead, Bruckert merely discloses routing in general. Moreover, Bruckert does not disclose a relative degree of overload. Instead, as already discussed herein, Bruckert is completely silent with respect to burden and load-balancing. Since Bruckert does not even remotely suggest a rate of routing or anything to do with overload, Bruckert cannot arguably be said to teach *adjusting the rate of routing based on a relative degree of overload on at least one of the plurality of servers*, as recited in claim 42.

Accordingly, Applicants respectfully request that the §103 rejection of claims 1-45 be withdrawn.

CONCLUSION

In view of the foregoing remarks, Applicants submit that all of the claims are patentably distinct from the prior art of record and are in condition for allowance. The Examiner is respectfully requested to pass the above application to issue. The Examiner is invited to contact the undersigned at the telephone number listed below, if needed. Applicants hereby make a written conditional petition for extension of time, if required. Please charge any deficiencies in fees and credit any overpayment of fees to Attorney's Deposit Account No. 09-0457.

Respectfully submitted,
Gordan G. GREENLEE et al.



Andrew M. Calderon
Registration No. 38,093

July 29, 2009
Greenblum & Bernstein, P.L.C.
1950 Roland Clarke Place
Reston, Virginia 20191
Telephone: 703-716-1191